

HIGH SPEED DATA STREAM PATTERN RECOGNITION

Related Applications:

This application is a non-provisional application of provisional application 60/322,012 filed 09/12/2001. Priority of the above referenced provisional is claimed.

Compact Disc Appendix:

This application includes a Compact disc appendix. The material on the compact disc is hereby incorporated herein by reference.

Field Of The Invention:

The present invention generally relates to systems and methods for performing, at high speeds, pattern recognition from streams of digital data.

Background Of The Invention:

With the continued proliferation of networked and distributed computers systems, and applications that run on those systems, comes an ever increasing flow and variety of message traffic between and among computer devices. As an example, the Internet and world wide web (the "Web") provide a global open access means for exchanging message traffic. Networked and/or distributed systems are comprised of a wide variety of communication links, network and application

1 servers, sub-networks, and internetworking elements, such as repeaters, switches,
2 bridges, routers, gateways.

3
4 Communications between and among devices occurs in accordance with defined
5 communication protocols understood by the communicating devices. Such
6 protocols may be proprietary or non-proprietary. Examples of non-proprietary
7 protocols include X.25 for packet switched data networks (PSDNs), TCP/IP for the
8 Internet, a manufacturing automation protocol (MAP), and a technical & office
9 protocol (TOP). Other proprietary protocols may be defined as well. For the most
10 part, messages are comprised of packets, containing a certain number of bytes of
11 information. The most common example is Internet Protocol (IP) packets, used
12 among various Web and Internet enabled devices.

13
14 A primary function of many network servers and other network devices (or nodes),
15 such as switches, gateways, routers, load balancers and so on, is to direct or
16 process messages as a function of content within the messages' packets. In a
17 simple, rigid form, a receiving node (e.g., a switch) knows exactly where in the
18 message (or its packets) to find a predetermined type of contents (e.g., IP address),
19 as a function of the protocol used. Typically, hardware such as switches and
20 routers are only able to perform their functions based on fixed position headers,
21 such as TCP or IP headers. Further, no deep packet examination is done.
22 Software, not capable of operating at wire speed is sometimes used for packet

1 payload examination. This software does not typically allow great flexibility in
2 specification of pattern matching and operates at speeds orders of magnitude
3 slower than wire rate. It is highly desirable to allow examination and recognition of
4 patterns both in packet header and payload described by regular expressions. For
5 example, such packet content may include address information or file type
6 information, either of which may be useful in determining how to direct or process
7 the message and/or its contents. The content may be described by a "regular
8 expression", i.e., a sequence of characters that often conform to certain expression
9 paradigms. As used herein, the term "regular expression" is to be interpreted
10 broadly, as is known in the art, and is not limited to any particular language or
11 operating system. Regular expressions may be better understood with reference to
12 Mastering Regular Expressions, J. E. F. Friedl, O'Reilly, Cambridge, 1997.

13
14 It is clear that the ability to match regular expressions would be useful for content
15 based routing. For this, a deterministic finite state automaton (DFA) or non-
16 deterministic finite state automaton (NFA) would be used. The approach used here
17 follows a DFA approach. A conventional DFA requires creation of a state machine
18 prior to its use on a data (or character) stream. Generally, the DFA processes an
19 input character stream sequentially and makes a state transition based on the
20 current character and current state. This is a brute-force, single byte at a time,
21 conventional approach. By definition, a DFA transition to a next state is unique,
22 based on current state and input character. For example, in prior art FIG. 1A, a

1 DFA state machine 100 is shown that implements a regular expression
2 "binky.*\jpg". DFA state machine 100 includes states 0 through 9, wherein the
3 occurrence of the characters 110 of the regular expression effect the iterative
4 transition from state to state through DFA state machine 100. The start state of the
5 DFA state machine is denoted by the double line circle having the state number "0".

6 An 'accepting' state indicating a successful match is denoted by the double line
7 circle having the state number "9". As an example, to transition from state 0 to
8 state 1, the character "b" must be found in the character stream. Given "b", to
9 transition from state 1 to state 2, the next character must be "i".

10
11 Not shown explicitly in FIG. 1A are transitions when the input character does not
12 match the character needed to transition to the next state. For example, if the DFA
13 gets to state 1 and the next character is an "x", then failure has occurred and
14 transition to a failure terminal state occurs. FIG. 1B shows part 150 of FIG. 1A
15 drawn with failure state transitions, wherein a failure state indicated by the "Fail"
16 state. In FIG. 1B, the tilde indicates "not". For example, the symbol "~b" means
17 the current character is "not b". Once in the failure state, all characters cause a
18 transition which returns to the failure state, in this case.

19
20 Once in the accepting state, i.e., the character stream is "binky.*\jpg", the receiver
21 node takes the next predetermined action. In this example, where the character
22 stream indicates a certain file type (e.g., ".jpg"), the next predetermined action may

1 be to send the corresponding file to a certain server, processor or system.

2
3 While such DFAs are useful, they are limited with respect to speed. The speed of a
4 conventional DFA is limited by the cycle time of memory used in its implementation.

5 For example, a device capable of processing the data stream from an OC-192
6 source must handle 10 billion bits/second (i.e., 10 gigabits per second (Gbps)).

7 This speed implies a byte must be processed every 0.8 nanosecond (nS), which
8 exceeds the limit of state of the art memory. For comparison, high speed SDRAM
9 chips implementing a conventional DFA operate with a 7.5 nS cycle time, which is
10 ten times slower than required for OC-192. In addition, more than a single memory
11 reference is typically needed, making these estimate optimistic. As a result,
12 messages or packets must be queued for processing, causing unavoidable delays.

13
14 **Summary Of The Invention:**

15 A system and method in accordance with the present invention determines in real-
16 time whether a set of characters from a data or character stream (collectively "data
17 stream") satisfies one or more of a set of predetermined regular expressions. A
18 regular expression may be written in any of a variety of codes or languages known
19 in the art, e.g., Perl, Python, Tcl, grep, awk, sed, egrep or POSIX expressions.

20 Additional means may be implemented to determine a next action from satisfaction
21 of one such regular expression, or from a lack of such satisfaction of a regular
22 expression. The present invention provides, an improved high speed, real-time
23 DFA, called a **Real-time Deterministic Finite state Automaton** (hereinafter RDFA).

1 The RDFA provides high speed parallel pattern recognition with relatively low
2 memory storage requirements. The RDFA includes a DFA optimized in accordance
3 with known techniques and a set of alphabet lookup and state related tables that
4 combine to speed up the processing of incoming data streams.

5
6 The data stream may be received by a typical computer and/or network device,
7 such as a personal computer, personal digital assistant (PDA), workstation,
8 telephone, cellular telephone, wireless e-mail device, pager, network enabled
9 appliance, server, hub, router, bridge, gateway, controller, switches, server load-
10 balancers, security devices, nodes, processors or the like. The data stream may be
11 received over any of a variety of one or more networks, such as the Internet,,
12 intranet, extranet, local area network (LAN), wide area network (WAN), telephone
13 network, cellular telephone network, and virtual private network (VPN).

14
15 An RDFA system in accordance with the present invention includes a RDFA
16 compiler subsystem and a RDFA evaluator subsystem. The RDFA compiler
17 generates a set of tables which are used by the RDFA evaluator to perform regular
18 expression matching on an incoming data stream. The data stream may present
19 characters in serial or parallel. For example, four characters at a time may arrive
20 simultaneously or the four characters may be streamed into a register. The RDFA
21 evaluator is capable of regular expression matching at high speed on these
22 characters presented in parallel.

1 The RDFA compiler subsystem generates a DFA state machine from a user
2 specified regular expression. The DFA state machine is optimized to include a
3 minimum number of states, in accordance with known techniques. A number of
4 bytes to be processed in parallel is defined, as M bytes. For each state in the state
5 machine, the RDFA compiler determines those characters, represented by bytes,
6 that cause the same transitions. Those characters that cause the same transitions
7 are grouped into a class. Therefore, each class, for a given current state of the
8 state machine, includes a set of characters that all cause the same transitions to
9 the same set of next states. Each class is represented by a class code. The
10 number of bits required for a class code is determined solely from the number of
11 classes at a given state and byte position.

12
13 The RDFA compiler generates a set of state dependent alphabet lookup tables
14 from the class codes and the relevant alphabet. A lookup table for a given state
15 associates a class code to each character in the relevant alphabet. Each of M
16 bytes under evaluation has its own lookup table. Classes are a compressed
17 representation of the alphabet used in a state machine, since multiple symbols can
18 be represented by a single class. This can lead to large reductions in the number
19 of bits required to represent alphabet symbols, which in turn leads to large
20 reductions in the size of next state lookup tables.

21
22 The RDFA compiler then generates a set of state dependent next state tables using

1 the class codes and state machine. For each state (as a current state) in the state
2 machine, a set of next states is determined and represented in a next state table.
3 For a given current state, the next state is a function of the characters represented
4 by the bytes under evaluation. The possible sets of concatenated class codes from
5 the alphabet lookup tables serve as indices to the possible next states in the
6 appropriate next state table. Given a current state, a table of pointers may be
7 defined, wherein each pointer points to the appropriate next state table from the set
8 of next state tables. The RDFA compiler can also determine the memory
9 requirements for RDFA system data associated with a defined regular expression.

11 During parallel evaluation, the RDFA evaluator selects, or accepts, the next M bytes
12 and gets the appropriate M lookup tables to be applied to the bytes under
13 evaluation. Each byte is looked up in its corresponding lookup table to determine
14 its class code. As previously mentioned, the class codes are concatenated. Given
15 a current state, the RDFA evaluator retrieves the appropriate next state table. The
16 code resulting from concatenation of the class code lookup results, is applied as an
17 index to the selected next state table to determine the next state which involves M
18 transitions beyond the current state.

20 This process continues until evaluation is terminated or the regular expression is
21 satisfied. The process may be terminated when, for example, the bytes under
22 evaluation do not cause a transition to a non-failure state. With a regular

1 expression satisfied, the next action may be determined by the RDFA system, or by
2 a system interfaced therewith. The RDFA system may be employed in any of a
3 variety of contexts where it is essential or desirable to determine satisfaction of a
4 regular expression, whether anchored or unanchored, in a data stream, particularly
5 when such determinations are to be made at high speeds, such as required by OC-
6 192 rates. The RDFA system may also be employed in contexts where
7 consumption of relatively small amounts of memory by the RDFA system data are
8 required or desirable.

9 10 **Brief Description Of The Drawings**

11 The foregoing and other objects of this invention, the various features thereof, as
12 well as the invention itself, may be more fully understood from the following
13 description, when read together with the accompanying drawings, described:
14

15 FIG. 1A is a state diagram implementing a regular expression, in accordance with
16 the prior art;

17 FIG. 1B is a portion of the state diagram of the regular expression of FIG. 1A,
18 including a failure state;

19 FIG. 2A is a block diagram of a RDFA system in accordance with the present
20 invention;

21 FIG. 2B is a block diagram of a RDFA compiler, from the RDFA system of FIG. 2A;

22 FIG. 2C is a block diagram of a RDFA evaluator, from the RDFA system of FIG. 2A;

1 FIG. 3 is a diagram depicting 4 byte parallel processing and 4 corresponding
2 alphabet lookup tables, used by the RDFA evaluator of FIG. 2C;
3 FIG. 4 is a diagram depicting a next state table, used by the RDFA evaluator of
4 FIG. 2C;
5 FIG 4A is a diagram indicating the flow of data from the character tables, the index
6 table and memory.
7 FIG. 5 is a diagram depicting characters that cause the same state transitions, used
8 by the RDFA compiler of FIG. 2B; and
9 FIG. 6 is a diagram depicting a state machine used by the RDFA compiler of FIG.
10 2B.
11 FIG. 6A illustrates a number of states reachable by 2-closure.
12 FIG 7A illustrates a DFA for processing 8 characters.
13 FIG 7B illustrates an RDFA for processing 4 bytes in parallel

14
15 **Detailed Description Of The Preferred Embodiments:**

16 In the preferred embodiment, the present invention is implemented as a RDFA
17 system 200, shown in FIG. 2A, which includes two subsystems. The first
18 subsystem is a RDFA compiler 210 that performs the basic computations
19 necessary to create tables for subsequent real-time pattern recognition. The
20 second subsystem is a RDFA evaluator 250 that performs the evaluation of
21 characters using the RDFA tables created by the RDFA compiler 210. The RDFA
22 system 200 includes a first memory 220 for high speed access by RDFA evaluator

1 250 during evaluation of characters from the data stream. This first memory 220
2 consists of on-chip or off-chip or any combination thereof. A second memory 204
3 includes the initial one or more regular expressions of interest, and need not lend
4 itself to high speed access, unless required as a function of a particular application
5 to which the RDFA is applied.

6
7 FIG. 2B is a block diagram of the RDFA compiler 210. As will be discussed in more
8 detail below, the RDFA compiler 210 includes a regular expression compiler 212
9 that converts a regular expression, from memory 204, into an optimized state
10 machine. An alphabet lookup table generator 214 generates, from the regular
11 expression and the state machine, a series of state dependent alphabet lookup
12 tables. The alphabet lookup tables include codes associated with each character in
13 an applicable alphabet of characters. These alphabet lookup tables are stored in
14 high speed memory 220. During RDFA data stream processing (i.e., character
15 evaluation), a character represented by a byte under evaluation is looked up in a
16 corresponding alphabet lookup table to determine its state dependent code, as will
17 be discussed in greater detail. A next state table generator 216 generates a table
18 of next states of the state machine to be applied during evaluation of a set of
19 characters, wherein next states are determined as a function of a current state and
20 the character codes from the alphabet lookup tables. The next state table is also
21 preferably stored in high speed memory 220.

FIG. 2C is a functional block diagram of the RDFA evaluator 250. The RDFA evaluator 250 includes several functional modules that utilize the alphabet lookup tables and next state tables generated by the RDFA compiler 210. At a top level, a byte selector module 252 captures the requisite number of bytes (i.e., M bytes) from an incoming data stream 205. An optional bit mask 251 can filter the input stream to select words from predetermined positions, allowing the processing to ignore certain portions of the input stream. Each bit in the mask corresponds to a four byte section of a packet for this embodiment. The selected bytes are taken and processed in parallel by an alphabet lookup module 254, which selectively applies the alphabet lookup tables from memory 220 to determine a character class code for each byte. As will be discussed in greater detail, characters causing the same state transition are grouped in classes, which are represented in alphabet lookup tables as class codes of a certain bit width. The alphabet lookup module 254 concatenates the class codes obtained from the lookup tables and passes the concatenated code to a next state module 256. The next state module 256 selectively applies the concatenated class codes to the appropriate next state table from memory 220, given a current state, to determine a next M th state in a corresponding state machine. This process continues at least until a failure state or accepting state is achieved.

The RDFA evaluator 250, as well as the RDFA compiler 210, may be implemented in hardware, software, firmware or some combination thereof. In the preferred

1 form, the RDFA evaluator 250 is a chip-based solution, wherein RDFA compiler 210
2 and high speed memory 220 may be implemented on chip 270. Memory 204 may
3 also be on-chip memory or it may be off-chip memory, since high-speed is typically
4 not as vital when generating the RDFA. However, if high-speed is required the
5 RDFA compiler 210 and memory 204 may each be on-chip. Therefore, preferably,
6 to achieve higher speeds the primary functionality of RDFA evaluator 250 for
7 processing incoming data streams is embodied in hardware. The use of pointers to
8 next state tables, rather than directly using the alphabet table lookup results, allows
9 flexibility in memory management. For example, if on-chip and off-chip memory is
10 available, then pointers can be used so that more frequently used memory is on-
11 chip, to speed up RDFA performance. The RDFA expression compiler 210 will
12 determine the amount of memory required. This allows the user to know if a
13 particular set of rules will fit in the on-chip memory. Thus, memory related
14 performance can be accurately known ahead of time.

15
16 As will be appreciated by those skilled in the art and discussed in further detail
17 below, a RDFA system 200 in accordance with the present invention requires
18 relatively modest amounts of high speed or on-chip memory 220, certainly within
19 the bounds of that which is readily available. Memory 220 is used to store the
20 alphabet lookup tables and next state tables for a given regular expression.

21 Unlike a conventional (i.e., single byte at a time processing) DFA approach, a
22 RDFA is configured for scalable parallel processing. As a general rule, increasing

1 the number of bytes (M) processed in parallel yields increasingly greater processing
2 speeds, subject to the limitations of other relevant devices. While in the preferred
3 embodiment provided herein, the RDFA evaluator 250 processes four (4) bytes in
4 parallel (i.e., $M = 4$), there is no inherent limitation to the number of bytes that can
5 be processed in parallel.

6 7 Data Stream Evaluation:

8 FIG. 3 illustrates the multiple alphabet lookup table concept 300 for a set of 4 bytes
9 320, which are selected from the data stream 205 by byte selector 252 and are
10 taken as parallel input 260 by alphabet lookup module 254 (see FIG. 2C). Each
11 byte represents a character (e.g., a number, a letter, or a symbol) from the
12 permitted alphabet. In the preferred embodiment, a separate alphabet lookup table
13 having 256 elements is defined for each of the 4 bytes and each state and is stored
14 in memory 220. The alphabet lookup tables 310 are formed and applied as a
15 function of a current state of a state machine that represents the regular
16 expression.

17
18 In the example of FIG. 3, a first alphabet lookup table 312, having a 2 bit width, is
19 used to lookup a first byte 322. A second alphabet lookup table 314, having a 3 bit
20 width, is used to lookup a second byte 324, and so forth with alphabet tables 316
21 and 318 and third byte 326 and fourth byte 328, respectively. The elements of the
22 alphabet lookup tables 310 are related to state transitions for a corresponding state

1 machine that models the regular expression. Accordingly, the selection and
2 application of alphabet lookup tables 310 is a function of the current state of the
3 state machine. The current state is the last state resulting from the processing of
4 the previous 4 characters, if any. Thus, a different set of alphabet lookup tables is
5 used for each current state.

6
7 The widths of the table entries for each byte can vary from one state to the next,
8 depending on the regular expression and the current state of the corresponding
9 state machine. In the FIG. 3 example, the table widths in bits are 2 bits for table
10 312, 3 bits for table 314, 3 bits for table 316, and 4 bits for table 318. The table
11 widths in another state might be 1 bit, 1 bit, 2 bits, and 4 bits, as an example. For
12 instance, if for the first byte there are only two possible character classes, then the
13 width of the alphabet lookup table for that bit need only be 1 bit. The current state
14 is stored in memory (e.g., on-chip memory 220) for use in determining which
15 alphabet lookup tables to apply to the 4 bytes 320 and for determining a next state.

16
17 For each of the 4 bytes 320, using lookup tables 310 a different class code is
18 obtained by alphabet lookup module 254. As previously discussed, the characters
19 are grouped into classes according to the state transitions the characters cause
20 and codes associated with those classes (i.e., class codes) are represented in the
21 alphabet lookup tables. Therefore, if byte 322 represents the character "a",
22 alphabet lookup module 254 finds the element in alphabet lookup table 312 that

1 corresponds to "a" and obtains the class code stored at that element (e.g., class
2 code 01). This is done for each other byte (i.e., bytes 324, 326 and 328) using their
3 respective alphabet lookup tables (i.e., tables 314, 316 and 318).

4
5 The lookup table class codes for each of the 4 bytes are concatenated together,
6 which for the FIG. 3 example produces a 12 bit result (i.e., 2 + 3 + 3 + 4 bits). As
7 an example, assume that from lookup tables 310 of FIG. 3 resulted a 2 bit word
8 "01" from table 312, a 3 bit word "001" from table 314, a 3 bit word "011" from table
9 316, and a 4 bit word "0000" from table 318. The resulting 12 bit concatenated
10 word would be "010010110000".

11
12 As is shown in FIG. 4, the current state of the state machine is used as an index
13 into a table of pointers 410. Table 410 is defined as a function of the regular
14 expression's state machine, so each current state has a corresponding table to
15 possible next states. Each pointer in table 410 points to a linear (i.e., 1 dimensional
16 (1-D)) table 420 of next state values (or a "next state table") and the 12 bit
17 concatenated result of the parallel alphabet lookup is used as an offset or index into
18 the selected next table 420. Therefore, a next state value is selected from next
19 state table 420 as a function of the current state and the concatenated 12 bit word.

20 The selected next state value corresponds to the next state. The next state
21 determined from evaluation of the 4 bytes serves as the current state for evaluation
22 of the next set of 4 bytes.

1
2 In the preferred form, the selected next state table value includes a terminal state
3 code (e.g., with higher order bit set to 1) that indicates whether or not the next state
4 is an accepting state (or terminal state). Generally, a terminal state is a state the
5 process enters when processing from a data stream with respect to a certain one or
6 more regular expressions is completed; i.e., it is indicative of termination of
7 processing with respect to the one or more regular expressions. For example, in
8 the preferred embodiment a high order bit associated with one or more of the bytes
9 under evaluation is set to "1" upon transition into a terminal state. In one
10 embodiment, the hardware stores the word (i.e., the 4 bytes under evaluation) for
11 which the terminal state occurred and the corresponding offset from the lookup
12 table (i.e., the 12 bit concatenated word). Thereafter, post processing software
13 may use the stored data to determine at which of the 4 bytes the regular expression
14 terminated. This is useful in many situations where only a small number of regular
15 expression matches occur per packet, so the number of such determinations is
16 relatively small. In another embodiment, the codes (i.e., the 4 bytes and 12 bit
17 word) are stored in a secondary terminal state table, which allows the hardware to
18 directly determine which byte terminated the processing. The benefit of allowing
19 the hardware to make such determinations is that it can be accomplished much
20 more quickly in hardware, which is a significant consideration in high speed, real-
21 time processing.
22

1 In accordance with the preferred embodiment, only three (3) memory operations
2 are required to process the 4 bytes. They are: (i) find characters in lookup tables
3 310; (ii) find pointer in table 410; and (iii) get next state indicia from next state table
4 420. Further, these operations may be easily pipelined by performing the character
5 table lookup at the same time as the last 4 byte result is being looked up in the next
6 state table. to allow improved processing times , with the only significant limitation
7 being the longest memory access.

8
9 The benefits of the preferred embodiment can be further appreciated when the
10 RDFA memory requirements are compared with those of a naïve DFA approach,
11 where the lookup is applied to a 4 byte word. In this type of DFA parallelization, 4
12 bytes would be looked up in parallel. This would require a table having 256^4
13 entries, which is about 4.295 billion entries, and a word (4 byte) cycle time of 3.2 nS
14 in order to keep up with OC-192 rates (i.e., 10 Gb/sec). This is impractical to
15 implement with current or near-term memory technology, based on the speed and
16 size required to keep up with OC-192 rates. Further, such a large amount of
17 memory cannot be implemented on-chip, so a significant amount of off-chip
18 memory would be required, unacceptably slowing the process. Compare the
19 memory requirement of simple DFA parallelization with the greatly reduced amount
20 of memory used in the preferred embodiment of the RDFA system 200. Note that
21 the naïve DFA parallelization requires many orders of magnitude greater memory
22 size than an RDFA system 200, in accordance with the present invention.

Fig. 4A is another illustration of the alphabet lookup tables 310, the index table 410, and the next state table 450 showing their operation and interactions. The bytes which are being examined are designated 320. Bytes 320 are a four byte segment from a data stream (i.e. four bytes from a data packet that is being examined). The alphabet lookup tables 310 have a segment associated with each possible state of the state machine. In Fig. 4A the states are designated s1, s2, s3, etc. along the left side of the figure. In each state, the bytes 320 are used to interrogate the section of table 310 associated with that particular state. The lookup operation produces a multi-bit result 330. The number of bits in the result 330 (i.e. the number of bits retrieved or generated by the alphabet lookup table) is a function of the particular bytes 320, the particular state, and the byte position. The index table 410 has an entry for each state. Each entry in table 410 includes a code which tells the system how many bits to retrieve from the output of the lookup table 310 for that particular state. (Note in an alternate embodiment this code is stored in a separate table that has a location for each state similar to table 410) During any particular state conventional addressing circuits address and read the contents of the location in table 410 associated with the particular state. The result bits 330 and the bits in the current state position of the index table 410 are concatenated to produce a memory address 441. As indicated at the left side of Fig. 4A, different locations in index table 410 have a different number of bits. The number of bits in concatenated result 441 (i.e. total number of bits from table 410 and result 330) is

1 always a fixed number. In the preferred embodiment that number is 13. The
2 number of bits equals the number of bits in a memory address for the particular
3 system. Thus, for each state the associated location in table 410 indicates how
4 many bits should be retrieved from table 310 and it provides a series of bits so that
5 there is a total of 13 bits for address 441.

6
7 Address 441 is the address of an entry in the next state table 450. The memory
8 address 441 is used to interrogate next state table 450 utilizing conventional
9 memory addressing circuitry. The entry in next state table 450 at address 441
10 indicates the next state. The entry in the next state table 450 may also contain a
11 flag which indicates that the operation has reached a special point such as a
12 termination point.

13
14 The operations proceed until the flag in the next state table indicates that the
15 operation has reached a termination point or that the bytes have been recognized
16 or matched. When a match is found, processing the bytes in a particular packet
17 can then either terminate or the system can be programmed to continue processing
18 other sets of bytes 320 in an attempt to find other matching patterns.

19
20 If the next state table does not indicate that the operation has terminated, the
21 process proceeds to the next state and the process repeats. If the process repeats
22 the information in appropriate next state table 450 is used. That is, the designation

1 of the next state in table 450 is used to generate the address of an appropriate
2 section of lookup table 310 and the process repeats. Upon reaching a termination
3 state, the following data is saved in memory registers 442:

- 4 1. Pointer to the word (4 bytes) in the packet at which the terminal state
5 occurred.
- 6 2. The table offset (computed from the alphabet table lookups results and index
7 table) into the next-state table.

8 The saved data can be used by post processing operations which determine what
9 action to take after the operation has terminated. In some embodiments when a
10 termination flag is encountered which indicates that a match is found, the operation
11 continues, that is, additional bytes in the string is are processed in an effort to
12 locate another match to the specified regular expression.

13
14 In general after four bytes have been processed, four different bytes are streamed
15 into register 320 and the process repeats. Furthermore, one can search for a wide
16 array of different patterns. A target pattern can be more than four bytes long. For
17 example if one is searching for a five byte pattern, after four of the bytes have been
18 located another set of four bytes can be streamed into register 320 to see if the fifth
19 byte is at an appropriate location.

20
21 An appendix on a CD is provided containing a specific example of the data that
22 would be stored in table 310, 410 and 450 so that the system would proceed

1 through a series of states to locate the character string "raqia". It is noted that each
2 different set of regular expressions which one wants to locate require a different set
3 of data in tables 310, 410 and 450. The example given is an example that contains
4 5 particular characters in sequence. It should however be understood that the
5 invention can be used to locate any desired regular expression, not just fixed
6 character sequences. The specific data for tables 310, 410 and 450 given in the
7 appendix are for locating or recognizing the particular character sequence "raqia".
8 The data files in the appendix are designated as follows: (a) the data for the four
9 byte positions for table 310 are designated: _hwct_0.txt, _hwct_1.txt,
10 _hwct_2.txt, _hwct_3.txt. (b) The data for index table 410 is designated _it.txt.
11 (c) The data for the next state table 450 is designated _nst.txt.

12
13 In the specific example provided in the appendix, the tables provide for 32 states of
14 operation. The four tables 310 each have 32 sections each with 256 entries for a
15 total of 8192 entries. The index table has 32 entries. It is noted that the choice of
16 32 states is matter of engineering choice for the particular application. In the
17 particular example given in the appendix, the next state table 450 has 8192 entries.
18 It is noted that the number of entries in this table is also a matter of choice. The
19 number of entries in the next state table for each state is determined by the number
20 of combinations of character classes for that state for all the byte positions. For
21 example, if the number of character classes for byte positions 0 through 3 are 4, 4,
22 8, 8 respectively, then the total number of next state table entries for that state is 4

1 x 4 x 8 x 8=1024. And the total size of the address space for all the states is the
2 sum of the table sizes for each state. In one embodiment the number of character
3 classes at each byte position is a power of 2, but other embodiments use various
4 different numbers of character classes.

5
6 It should be noted that while in the embodiment described, four bytes are
7 processed in parallel, alternate embodiments can be designed to handle different
8 numbers of bits in parallel. For example other embodiments can handle 1, 2, 6, 8,
9 12 bytes in parallel.

10 11 Creation of the RDFA Tables:

12 To generate a RDFA in accordance with the present invention, the regular
13 expression compiler 212 converts a regular expression from memory 204 into a
14 DFA. The regular expression compiler 212 may also optimize the DFA to minimize
15 the number of states. These processes are known in the art, so are not discussed
16 in detail herein. The regular expression compiler is also configured to determine
17 the amount of memory required to store the RDFA for a given regular expression,
18 as will be discussed in further detail below. This allows the user to know if a
19 particular set of rules (i.e., regular expressions) will fit in the on-chip memory. Thus,
20 performance can be accurately predicted.

21
22 The regular expression compiler 212 also reduces state transition redundancy in

1 the alphabet representing the input data stream by recognizing that DFA state to
2 state transition decisions can be simplified by grouping characters of an alphabet
3 according to the transitions they cause. The list of states that may be reached in a
4 single transition is referred to as '1-closure'. The term " n -closure" is defined as the
5 list of states reachable in n transitions from the current state. n -closure is readily
6 calculated recursively as the list of states reachable from the $n-1$ closure. There
7 may be more than one character that causes the same transitions to the same n -
8 closure set. In such a case, characters may be grouped into classes according to
9 the set of next state transitions they cause. Rather than representing individual
10 characters, each class may be represented in a 1, 2, 3, or 4 bit code, for example.
11 In this manner, the applicable alphabet is represented in an extremely economical
12 form.

13
14 Even very complicated expressions can achieve significant compression in the
15 number of bits required to represent its alphabet by mapping to character classes.
16 For example, a portion of a regular expression represented as "(a | b | c | g)" can be
17 represented in a state transition diagram 500, shown in FIG. 5, wherein the
18 expression indicates "a" or "b" or "c" or "g". These characters all cause a transition
19 from state "1" to state "2", therefore, these characters can all be mapped into a
20 single class. If all other characters cause a transition to a failure state, then all of
21 those characters can be grouped into a second class. Therefore, when in state 1 of
22 FIG. 5 (i.e., state 1 is the current state) all transitions can be represented with a 1

1 bit code, wherein a code of "1" could indicate a transition into state "2" and a code
2 of "0" could indicate a transition into a failure state F (not shown). Mapping
3 characters into classes eliminates the need to represent characters with 8 bits, as is
4 typical with conventional approaches.

5
6 Alphabet lookup tables are generated by the alphabet table generator 214 of FIG.
7 2B. In the present invention, the RDFA alphabet lookup tables reflect the classes
8 that represent the state transitions, wherein the 1, 2, 3, or 4 bit representation, as
9 the case may be, are embodied in character classes related to the current state. In
10 general, when M bytes are processed in parallel a separate alphabet lookup table is
11 computed for each of the M bytes. Further, a different set of tables is computed for
12 each state in the state machine. Thus, if a state machine has L states and M bytes
13 are processed in parallel, a total of $(L \times M)$ alphabet lookup tables are produced, in
14 accordance with the preferred embodiment.

15
16 The algorithm used to produce the M character class tables for a regular
17 expression state machine from a starting state S , is as follows. The n th alphabet
18 lookup table (where $1 \leq n \leq M$) uses the previously computed $n-1$ closure and then
19 computes the n -closure. Then, for each character in the alphabet, a list of non-
20 failure state transitions from the $n-1$ closure to the n -closure is generated. An
21 alphabet mapping is then initialized by placing the first character in the alphabet
22 into character class 0. The transition list for the next letter in the alphabet, for a

1 given regular expression, is examined and compared with the transitions for the
2 character class 0. If they are identical, then the character is mapped to class 0,
3 otherwise a new class called "class 1" is created and the character is mapped to it.
4 This process proceeds for each character in the alphabet. So, if a list of transitions
5 for a character matches the transitions for an existing class, then that character is
6 represented in that existing class, otherwise that character is the first member of a
7 new class. The result of this process is a character class number associated with
8 each character in the alphabet. The total number of classes for a particular lookup
9 table may be represented by P . Then, the number of bits necessary to represent
10 each symbol is given by:

$$Q = \text{floor}(\log_2 P) + 1$$

12 Q is also the width of the table entries in the alphabet lookup table (e.g., 1, 2, 3, or
13 4 bits). For example, in alphabet lookup table 312 of FIG. 3, $Q = 2$. Note that Q is
14 computed for each alphabet lookup table separately and varies as a function of
15 both state and byte position.

17 This concept may be appreciated with a simple example for processing 2 bytes in
18 parallel (i.e., for $M = 2$) for the portion 650 of a state machine 600 shown in FIG. 6.

19 This example focuses primarily on lookup tables and transitions from the 0 state.

20 State machine 600 is derived from a predefined regular expression. The 1-closure
21 for state 0 is (1, 2, 3, F), where the failure (or terminal) state may be denoted by
22 symbol F (not shown). That is, as can be seen from FIG. 6, from state 0, non-

failure transitions may be to state 1, state 2, or state 3. Table 1 provides a list of state transitions out of state 0 for an alphabet consisting of letters from "a" to "k", in accordance with the state diagram 600 of FIG. 6.

Letter	Transitions
a	$0 \Rightarrow 1$
b	$0 \Rightarrow 1$
c	$0 \Rightarrow 2$
d	$0 \Rightarrow 2$
e	$0 \Rightarrow 3$
f	$0 \Rightarrow 3$
g	$0 \Rightarrow 3$
h	\emptyset null (or F state)
i	\emptyset null (or F state)
j	\emptyset null (or F state)
k	\emptyset null (or F state)

Table 1 - Transitions for State Diagram 600, From State 0

Upon inspection, Table 1 shows that the alphabet maps to 4 different equivalent classes, meaning that 2 bits are sufficient for the width of an alphabet lookup table for a current state of state 0. Therefore, with regard to a current state 0, the following classes may be formed: class 0 (a, b), class 1 (c, d), class 2 (e, f, g,) and a failure state class 3 (h, i, j, k). In the corresponding alphabet lookup table, class 0

may be represented as "00", class 1 as "01", class 2 as "10", and class 3 as "11",
as follows:

Letter	Class Code In Alphabet Lookup Table
a	00
b	00
c	01
d	01
e	10
f	10
g	10
h	11
i	11
j	11
k	11

Table 2 - Lookup Table Entries, Current State 0

The 2-closure for state machine 600 is (1, 4, 5, 6, 7, 8, *F*) from state 0. Similarly,
Table 3 is a list of state transitions for each character for the 2-closure. In this
case, inspection of Table 3 shows the alphabet maps to 8 equivalent characters, so
that 3 bits are required for the table width. Note that as indicated in the *Q* value
calculation, if the number of equivalent characters had been 5, the table width
would still be 3 bits.

Letter	Transitions
a	$1 \Rightarrow 4$
b	\emptyset null (or F state)
c	$1 \Rightarrow 1$
d	$1 \Rightarrow 4$ $2 \Rightarrow 5$
e	$3 \Rightarrow 6$
f	\emptyset null (or F state)
g	\emptyset null (or F state)
h	$2 \Rightarrow 6$
i	$3 \Rightarrow 7$
j	$3 \Rightarrow 8$
k	\emptyset null (or F state)

Table 3 - Transitions for State Diagram 600, From State 1

The next state table generator 216 of FIG. 2B generates, for each state in the state machine (as a current state), a list of next states represented in next state tables (e.g., next state table 420 of FIG. 4). Therefore, for each current state, there is a one-dimensional (1-D) list of possible next states. Using state machine 600 of FIG. 6, assume the current state is state 0 and $M = 4$ (where M is the number of bytes processed in parallel). As was discussed with respect to states 0 through 3 above, and shown in Tables 1 - 3, the characters that cause state transitions can be

1 mapped into classes. The corresponding next state table is comprised of entries
2 that dictate the next state given the class codes of the four characters (or bytes)
3 under evaluation from the alphabet lookup tables.

4
5 Assume 4 bytes were received representing the 4 characters "c, h, i, e". As
6 mentioned previously, the class code for "c" with a current state 0 is 01. The class
7 code for "h" is 3 bits, as the second of 4 bytes assume its class code is 011. Also,
8 assume for a current state 0 and as the third of 4 bytes, the class code for "i" is
9 0011. Finally, assume for a current state of 0 and as the fourth byte, the class code
10 for "e" is 101. The corresponding next state table, will have a next state value
11 corresponding to state 12, given the above class codes for the 4 bytes and a
12 current state 0. In the preferred form, the class codes are concatenated (e.g.,
13 010110011101) to form an index into the next state table, thus yielding the proper
14 next state. In this manner, the next state table and the corresponding table of
15 pointers, which are addressed by state, is generated for a regular expression. That
16 is, next state table generator 216 works through the state machine and alphabet
17 lookup tables to generate the next state and pointer tables.

18
19 The following is an explanation of the invention from a somewhat different
20 perspective: The purpose of the invention is high speed recognition of patterns in a
21 data stream. The patterns are described by 'regular expressions', which means
22 they may be quite general. For example, the regular expression to detect filenames

1 prefixed by 'binky' or 'winky', containing 'xyz' and having a filename extension '.jpg'
2 are found by the regular expression:

3 (binky | winky).*xyz.*\.

4 The RDFA (i.e. the present invention) can search for patterns at fixed locations
5 (anchored), as needed for IP packet filtering, but it can also locate unanchored
6 regular expressions anywhere in a packet payload.

7
8 The RDFA has other very important features and advantages over a conventional
9 DFA. It allows parallel processing of bytes. This is important in high speed
10 applications such as OC-192 transport layers, where four bytes arrive from the
11 framer at time. A conventional DFA can not be easily implemented at OC-192 rates
12 with todays memory speed cycle time and logic delay time limitations.

13
14 Another advantage is that the RDFA has memory requirements that can be
15 precomputed for a particular set of patterns to be recognized. Finally, the design
16 allows convenient separation of the algorithm between on and off-chip memory
17 when expression complexity becomes large.

18
19 A conventional DFA requires creation of a state machine prior to its use on a data
20 stream and the RDFA has a similar requirement. The user makes a list of regular
21 expressions (rules) and actions to be carried out if a rule is satisfied. A special
22 purpose RDFA compiler converts the rules and actions into a set of tables which

are downloaded to the RDFA hardware. The operation of the RDFA is described below, followed by a description of the algorithm implemented in hardware. Finally, the process and algorithms used to created the RDFA tables are described.

The RDFA operates in a manner similar to a (DFA) but with new features that allow parallelization of the processing, while making enormous reduction in the memory requirements compared with naïve parallelization. The speed of a conventional DFA is often limited by the cycle time of memory used in its implementation. For example, processing the data stream from an OC-192 source must handle 10 billion bits/second (10 Gbs). This speed implies a byte must be processed every 0.8 nS, which is beyond the limit of state of the art memory and logic circuits. For comparison, conventional high speed SDRAM chips operate with a 7.5 nS cycle time, which is ten times slower than required. A feature of the RDFA is the ability to process the data stream bytes in parallel. As described below the RDFA processes 4 bytes in parallel, but the algorithm may be applied to arbitrary numbers of bytes, meaning that it is scalable to higher speed. This allows the use of memory that is readily available and far lower in cost than required by a brute-force, single byte at a time, conventional approach.

In the specific embodiment presented, a separate 256 element lookup table is used for each of the 4 bytes processed in parallel, with table entries having 1 to 4 bits width per entry. Fig 3 illustrates the multiple alphabet lookup table concept. In this

1 particular example, the first byte looked up produces a two bit result. The second
2 byte produces a three bit result and so forth. It is very important to note that the
3 set of alphabet lookup tables used are state-dependent. Thus, a different set of
4 four tables is used for each state. Further, the widths of the table entries for each
5 byte position can vary from one state to the next. Thus, the width of each alphabet
6 table lookup is stored as a function of state, so that it can be looked up by state, in
7 order to concatenate the correct number of bits from each lookup. In the Fig 3
8 example the table widths in bits are denoted (2, 3, 3, 4). The table widths in
9 another state might be (1, 1, 2, 4).

11 The lookup results for each byte are concatenated together, which for the Fig 3
12 example produces a 12 bit result. Next, as shown in Fig 4, the current state is
13 used as an index into a table of pointers and a single pointer is selected on the
14 basis of current state. Each pointer, points at a separate linear (1-D) table of next
15 state values and the concatenated result of the parallel alphabet lookup is used as
16 an offset into the selected table. The selected table entry is the value of the next
17 state and also has a code indicating when a terminal state has been reached. The
18 high order bit for each next-state table entry is called the 'special flag' and is set to
19 one to indicate an acceptor state has been reached somewhere in the 4 bytes
20 being processed. In the preferred embodiment, when a lookup in the next-state
21 table results in the special flag being set, the hardware will store two pieces of
22 information into the next entry of a DFA results table. The data saved are:

1 1. Pointer to the word (4 bytes) in the packet at which the terminal state
2 occurred.

3 2. The table offset (computed from the alphabet table lookups results and
4 index table) into the next-state table.

5
6 Post processing software can use the saved data to determine at which of the 4
7 bytes the regular expression terminated. This is acceptable in situations of interest,
8 only a small number of regular expression matches occur per packet. Alternate
9 embodiments store codes into a secondary terminal state table which will let the
10 hardware directly determine which of the 4 bytes in a word, terminated the matched
11 pattern.

12
13 Regardless of how the special flag is set, the value of the next-state table entry is
14 used to set the next state of the RDFA machine. Thus, when processing the full
15 packet has finished, the entries in the result-table contain information that can be
16 used to determine all regular expression matches in the packet.

17
18 The use of pointers to next state tables, rather than directly using the alphabet table
19 lookup results, allows flexibility in memory management. For example, in
20 embodiments that have on-chip and off-chip memory, pointers can be used so that
21 more frequently used memory is on-chip, to speed up RDFA performance. The
22 expression compiler can determine the amount of memory required. This allows

1 the user to know if a particular set of rules will fit in the on-chip memory. Thus,
2 memory related performance can be accurately known ahead of time.

3
4 The preferred embodiment requires memory lookup operations to process the 4
5 bytes. Specifically, the memory lookups are:

- 6 1. Parallel lookup of each incoming byte from the data stream.
- 7 2. Lookup of the number of bits width for each alphabet lookup result based on
8 state
- 9 3. Lookup of pointer for next-state table based on current state
- 10 4. Lookup of next state.

11 These memory operations may be pipelined to allow effective processing times
12 limited by the longest memory access. Another advantage of the approach is seen
13 when its memory requirements are compared with a simple DFA approach applied
14 to processing 4 bytes in parallel. A simple approach to DFA parallelization, does a
15 lookup on the 4 bytes in parallel. This will match the speed of the RDFA, but
16 requires a table of size 2^{32} entries, which has 4.295 billion entries and a cycle time
17 of 3.2 nS in order to keep up with OC-192 rates (10 Gb/sec). Such a system is
18 difficult to implement with current or near-term memory technology, based on the
19 speed and size required. Further, such a large memory is difficult to implement on-
20 chip with the RDFA processing algorithm.

21
22 RDFA Algorithms and Programming (That is, Creation of the RDFA Tables): The

1 front end to the RDFA converts a regular expression to a Nondeterministic Finite
2 Automata (NFA). The NFA is then converted to a Deterministic Finite Automata
3 (DFA). Finally, an optimization is done on the DFA to minimize the number of
4 states. Such processes are well known to compiler writers and are described in
5 the literature. The process described herein converts the state-optimized DFA to
6 an RDFA.

7
8 The RDFA reduces the redundancy in the alphabet representing the input stream
9 by recognizing that at a given state in the DFA the transition decision can be made
10 by grouping letters of the alphabet according to the transition they cause. For
11 example, in the Fig 1 the only symbol which will cause a transition to state 3 is the
12 letter 'n'. From state 2, all other characters transition to the failure state. Thus, for
13 this case, a single bit is sufficient to represent the alphabet because the characters
14 can be mapped to two classes, 'n' and everything else. More complicated
15 transitions can still in general achieve large compression in the number of bits
16 required to represent the alphabet through mapping to character classes. For
17 example, a portion of a regular expression represented as: (a | b | c | g) would be
18 represented in a state transition diagram as shown in Fig. 5. The use of mapping
19 characters to classes changes the 8 bits conventionally used to represent the input
20 character to just a few bits. The basic idea behind mapping characters to classes is
21 recognition that as far as transitions in the state machine are concerned many
22 different characters have identical results. Thus, these many different character

1 codes, which have identical results as far as state transition can be mapped to the
2 same character code, referred to as a 'class'. The alphabet lookup tables perform
3 this mapping function at wire speed. Offline work is required to initially determine
4 them.

5
6 Algorithm for Production of Alphabet Lookup Tables: The list of states that may be
7 reached in a single transition is referred to as '1-closure'. We define the term 'n-
8 closure' as the list of states reachable in n transitions from the current state. n-
9 closure is readily calculated recursively as the list of states reachable from n-1
10 closure. In general when M bytes are processed in parallel, a separate alphabet
11 lookup table must be computed for each of the M bytes. Further a different set of
12 tables is computed for each DFA state. Thus if a DFA has L states and M bytes are
13 processed in parallel, a total of $(L \times M)$ alphabet lookup tables must be produced.

14 The algorithm used to produce the M character class tables used to parallel
15 process M bytes, for starting state S, is as follows. The n^{th} lookup table ($1 \leq n \leq M$)
16 production first computes the n-1 closure. For each symbol in the alphabet, a list of
17 state transitions from the n-1 closure to the next state is made. Now the alphabet
18 mapping is initialized by placing the first symbol in the alphabet into character class
19 0. The transition list for the next letter in the alphabet is examined and compared
20 with the transitions for the character class 0. If they are identical, then this
21 character is mapped to class 0, otherwise a new class called 'class 1' is created
22 and the character is mapped to it. This process proceeds for each character in the

1 character set. If the list of transitions for the character matches an existing class,
2 then it is placed into that class, otherwise it is the first member of a new class. The
3 result of this process is a character class number associated with each symbol in
4 the alphabet. This mapping means that all characters in a given class have an
5 identical list of transitions from the n-1 closure to the next state and thus as far as
6 the state machine behavior is concerned are identical symbols.

7
8 If the total number of classes for a particular lookup table is P . Then, the number of
9 bits, Q , necessary to represent each symbol is $\text{floor}(\log_2 P) + 1$. Q is the width of
10 the table entries in the lookup table. Note that Q is computed for each lookup table
11 separately and varies as a function of both state and byte position. Since the
12 above discussion is abstract, the concept is illustrated below with a simple example
13 for processing 2 bytes in parallel for the portion of a DFA shown in Fig. 6A. The
14 illustration will be done only for the lookup tables from the 0 state. The 1-closure
15 for this example is, (1, 2, 3, F) where the failure state is denoted by symbol F . The
16 2-closure is (1, 4, 5, 6, 7, 8, F) in this illustration. Table 1, which is given in an
17 earlier section of this document is a list of non-failure transitions out of state 0 for
18 an alphabet consisting of letters from a to k. Upon inspection, Table 1 shows that
19 the alphabet maps to 4 different equivalent characters, meaning that 2 bits are
20 sufficient for the table width. Similarly, Table 2 (also given in an earlier section of
21 this document) is list of non-failure state transitions for each symbol for 2-closure.
22 In this case the alphabet maps to 8 equivalent characters, so that 3 bits are

1 required for the table width. Note that as indicated in the Q value calculation, if the
2 number of equivalent characters had been 5, the table width, would still be 3 bits.

3
4 An important Feature of RDFA: An important property of the RDFA is that the bytes
5 in the data stream are treated as letters in an alphabet and are mapped to
6 character classes. In general, many characters map to a single class, greatly
7 reducing the number of bits necessary to represent an alphabet symbol. As a
8 consequence, when multiple characters are concatenated together and used for a
9 next-state table lookup, the size of the next-state table is greatly reduced, when
10 compared with concatenation of multiple bytes.

11
12 Important Hardware Implementation Feature: The RDFA has many applications,
13 some involving searching full packets for unanchored expressions. The system (i.e.
14 the engine) described above, is well suited to this application. Another application
15 is searching fixed headers for patterns. A special feature incorporated into the
16 RDFA is a programmable datastream bit mask, where each bit corresponds to a
17 sequential word in the input data stream of a packet. For example, an ethernet
18 packet containing 1500 bytes contains 375 words, and a 375 bit mask allows
19 complete freedom in selection of words to be processed. When a bit is set on in
20 the data stream mask, the corresponding word is fed to the RDFA. If the bit is
21 turned off then the corresponding word is not seen by the RDFA. This allows a
22 front end filter that operates at line rate which greatly reduces the load on the RDFA

1 when processing fixed position header information. Further, this can lead to
2 reductions in the complexity and memory used by the RDFA. With the above
3 described mask only a small subset of the datastream must be processed and the
4 data that is processed can be handled in a simpler manner, which in turn means
5 larger rule sets can be used for a given amount of memory.

6
7 Reduction of Table Sizes: The RDFA requires a set of alphabet lookup up tables
8 and a next state table for each state. If the number of states can be reduced, then
9 the size of the lookup tables can be reduced. In a classic DFA, when M characters
10 are processed the state machine transitions through M states. For an RDFA it is
11 recognized that processing M bytes in parallel can be treated as a black box,
12 transitioning between two states. For example, as shown in Fig 7A, the character
13 string 'abcdefgh' is intended to be matched. Not counting initial state, a classic
14 DFA has 8 internal states through which it transitions including the acceptor state.
15 However, if 4 bytes are processed in parallel, then only 2 states are needed to
16 represent the transitions as shown in Fig 7B. Note that this is a special case since
17 cyclic graphs, representing wild-cards or arbitrary numbers of character repetitions,
18 may not occur in this type of processing.

19
20 The invention may be embodied in other specific forms without departing from the
21 spirit or central characteristics thereof. While not discussed in detail, incoming data
22 may be evaluated against a plurality of regular expressions simultaneously. In such

1 a case, entering a failure state for one regular expression state machine only
2 terminates processing with respect to that regular expression. The present
3 invention may also be implemented in any of a variety of systems, e.g., to detect a
4 computer virus in e-mail. The present embodiments are therefore to be considered
5 in all respects as illustrative and not restrictive, the scope of the invention being
6 indicated by appending claims rather than by the foregoing description, and all
7 changes that come within the meaning and range of equivalency of the claims are
8 therefore intended to be embraced therein.

9 What is claimed is:
10